REMARKS

Reconsideration and allowance of the above-identified application are respectfully requested. The specification has been amended at lines 19-24 on page 7 to correct various errors. The abstract has been amended to reduce the number of words and more clearly summarize the specification. Claims 1-16 are currently pending in the present application. Claims 1, 2, 9 and 10 have been amended. No new matter has been added.

In the Official Action, the Abstract has been objected to. By way of the foregoing amendments, the Abstract has been amended to address the stated objections.

Claims 1-16 stand rejected under 35 U.S.C. §112, second paragraph, as allegedly being indefinite for failing to particularly point out and distinctly claim the subject matter set forth therein. By way of the foregoing amendments, each of the objections raised by the Examiner in numbered paragraph 6, subsections a, b and d have been addressed by removing the language which was alleged to be indefinite, although such cancellation of the language is not meant to indicate that the undersigned agrees that such language is indefinite. Regarding numbered paragraph 6, subsection c, it is respectfully submitted that claim language can not be considered indefinite merely because the specification does not "explicitly" define a record portion. This portion of the claim language merely sets up a relationship between a record portion and a data packet from which it is generated and, while broad, is not indefinite. Accordingly, reconsideration and withdrawal of this ground of rejection are respectfully requested.

Claims 1-16 stand rejected under 35 U.S.C. §103(a) as allegedly being unpatentable over Van Asten et al. (US 2003/0174699) in view of Rege et al. (U.S. Patent No. 5,390,299). This ground of rejection is respectfully traversed for at least the following reasons.

As regards the method of queuing variable size data packets according to Applicant's claim 1 combination, this claim set forth, among other things, generating from each packet a record portion of predetermined fixed size and containing

information about the packet. Data portions of the packets are stored in independent memory locations in a first memory while the record portions are stored in one or more managed queues in a second memory, whose memory locations are fixed and equal in size to the size of the record portions. The first memory is larger than the second. Moreover, and significantly, the memory locations in the first memory are in blocks having a plurality of different sizes and these memory locations are allocated to the data portions according to the size of the data portions.

In [¶0024] of the Van Asten publication, variable size packets are divided into a plurality (X) of segments (cells) all of which are of a single fixed-size. These segments/cells are stored in N memory modules such that the X cells are grouped into X/N groups of cells. In [¶0025] the cells of a group are stored in a selected different one of the N memory modules at an address selected from a subset of addresses according to a first relationship. In [¶0026], a multi-cell pointer (MCP) is stored for each group of cells. Each MCP has N memory module identifiers recording the order in which the cells of each group are stored in each of the selected N memory modules. In [¶0027] the MCP is stored at an address related to one of the addresses in the subset according to a second relationship. In the preferred embodiment, cells are assigned to the X/N groups of cells in sequence. In [¶0030] the preferred embodiment has N=5 whereas X ranges from 1 to about 200. In the method set out from [¶0046] to [¶0057], the MCP is created for each cell group and the memory module identifier of the appropriate one of the N memory modules is stored in the MCP prior to the MCP being stored in the MCP storage.

In practice, as described in [¶0096] of Van Asten, incoming variable size packets arriving at the switch of Figure 3 are divided into fixed size packets (ie cells or switch cells), each containing a payload from the original packet and a header, as illustrated in Figure 4. In [¶0097] large packets are segmented into a plurality of cells whilst in [¶0098] a very short packet is segmented into a single cell. The headers of the cells contain information to steer the cell through the switch fabric in addition to system-related information [¶0099]. In [¶0101] cells of a given segmented packet emerge from the switch fabric in the same order as their arrival time at the switch fabric, i.e., the cell order is preserved through the switch fabric.

The high speed packet memory of Figure 5 incorporates MCP storage and a number N of DRAM memory modules for the cells, as illustrated in Figure 6. In [¶0119] the MCP storage has the capacity to store up to a number M of MCPs and in [¶0120] each DRAM can store the same number M of switch cells. In [¶0121] and [¶0122] the MCPs and the cells forming a group are stored at the same address (A) in the MCP storage and in the DRAM modules. From Figure 6 in collaboration with Figure 4, and especially the manner in which the cell is identified by reference numberal 214, it appears that the DRAM address locations are all of the same size as one another, and are such that one cell can be stored in one location. In [¶0123] the MCPs indicate the location in the DRAMs of the first and subsequent cells of a given group, ie b1, b2…bN.

Figure 7 illustrates the MCP memory/queue organisation used to track the segmented packets in the high speed memory. Paragraph [¶0144] indicates that memory addresses are stored in a pool ("free list") of available addresses. In [¶0152] - [¶0154] as cells from the same packet arrive, the first is allocated an address in one of the DRAM modules and subsequent cells of the same segmented packet are allocated the same address in successive DRAM modules. The corresponding MCP in the same address in the MCP storage is updated each time with the relevant DRAM module number(s) for the subsequent cells. If a segmented packet contains more than N switch cells, additional MCPs and DRAM locations are assigned. The number of addresses needed to store such a packet equals the number of cells divided by M (N?) then rounded up. In the given example, where N=5, a large packet segmented into 26 cells would need 6 MCPs and addresses (ie 26/5 = 5.2 rounded up to 6) to store the switch cells from that packet.

However, it is respectfully submitted that there are some significant features taught in Van Asten that distance it from the claim 1 combination. First, Van Asten teaches that each cell segmented from a packet contains a payload and a header. In contrast, in claim 1, a record portion is generated and that record portion (and only the record portion) is stored in one or more managed queues, while the data portion (and only the data portion) is stored in the first memory. This is clearly not the case in Van Asten.

Second, Van Asten teaches providing as many extra memory locations as are necessary to accommodate larger packets. In Van Asten, all of the individual memory locations are of the same size [see, e.g., Figures 4 and 6]. Packet segments (cells) are all of the same size as one another [¶0096]. The individual cells of a given packet are distributed over as many DRAM locations as necessary. If a segmented packet cannot be stored at a single address across the N DRAMs, a further address is generated (ie taken from the free list) and the remaining cells are stored at another address across the DRAMs, the corresponding MCP being stored at the same address in the MCP store). There is no possibility that Van Asten can operate by allocating memory portions to the data portions according to the size of the data portions, as set forth in Applicant's claim 1 combination, where the individual memory locations for the data portions are of a plurality of different sizes.

As described in Applicant's specification, these sizes are predetermined (at least for any given data packet stream and not necessarily globally) to match the most likely packet sizes. In this way, the data portions of the variable size incoming packets can each be stored in a respective single memory location, a "small" packet being stored in a location of a first size and a "large" packet being stored in a location of a second size, the second size being larger than the first size. A significant advantage that this approach presents over Van Asten is that fewer segments need to be created from each data packet, thereby enhancing speed of operation. Also, the fact of having "small" memory blocks available for packet storage allows more efficient use of memory when a packet fits entirely into a single block. However, the provision also of "large" blocks means that when a packet does need to be split into portions, the resultant linked lists can themselves be managed more efficiently because accesses can be pipelined, with a consequence being that the overhead of the pointers to the next blocks is less significant.

By way of contrast, using a single cell size as in Van Asten means either that memory is wasted or even greater overhead is added for using multiple cells. Moreover, the fact that Van Asten creates cell segments, each of which contains a header portion, means that extra storage space and associated control are needed to store any given packet as compared to Applicant's solution. The fact that each

cell carries its own header binds the whole packet processing approach to the same constraints that the present invention is designed to overcome. There is, consequently, in Van Asten a teaching away from Applicant's claim 1 combination rather than any teaching which would have motivated one of ordinary skill in the art to have arrived at Applicant's claim 1 combination.

The Official Action correctly recognizes that Van Asten does not expressly disclose a second memory having a fixed memory size. However, as described above, it is respectfully submitted that the differences between Applicant's claim 1 combination run much deeper, given the two additional features mentioned above that are not taught in Van Asten.

The Official Action also acknowledges that Van Asten fails to expressly disclose a first memory which is larger than a second memory as set forth, among other features in Applicant's claim 1 combination. Overall, Van Asten fails to disclose at least four of the features Applicant's claim 1 combination, such that it seems unlikely that one of ordinary skill in the art would have been motivated to have arrived at the claimed combination based upon the disclosure of Van Asten whether taken singly or in combination with the Rege patent. In this latter regard, it is noted that, at least, Rege fails to remedy the above-described deficiencies of Van Asten.

Moreover, as stated in [¶0031] of the instant specification, a significant feature of the exemplary embodiments is that the record and data portions are handled separately and independently of one another. On the contrary, Van Asten creates packet segments (cells) each of which has a header and a payload portion. As a result, the Van Asten data portions are coupled with header portions and there is no decoupling of the data portions and record portions as achieved by these exemplary embodiments. Van Asten therefore suffers from the disadvantages outlined in eg [¶0007] and [¶0034] of the instant specification.

Similar comments as made above with respect to claim 1 are also applicable to independent claim 9. In order to make these distinctions clearer, claims 1 and 9 have been amended to emphasize that the first memory stores only data portions of the packets and the second memory stores only the record portions.

The dependent claims are allowable for at least the reasons set forth above

with respect to the independent claim from which they depend. Moreover, these claims are also allowable for reasons of their own.

For example, with regard to dependent claim 2, the Official Action refers to [¶0007] and [¶0008] and states that "each port card contains a separate packet memory which corresponds to two sizes of memory location in the first memory arranged in said two blocks". However, the mere fact that each port card contains a separate packet memory does not lead to the conclusion that there are two sizes of memory locations in the first memory. Rather, the memory locations in Van Asten are the same size from one port card to the other.

With regard to dependent claim 3, the Official Action draws a comparison between the fixed size packets (cells) in [¶0013] of Van Asten and the most commonly occurring sizes of data packet. However, it is respectfully submitted that Van Asten does not suggest in [¶0013] that the cell size is matched to the most commonly occurring packet sizes but suggests that the cell size is related to the bus capacity. Although there may ultimately be a connection between bus capacity and cell size, the claim 3 specifies a plurality of different sizes and this is not taught or suggested in Van Asten.

All of the objections and rejections raised in the Official Action having been addressed, it is respectfully submitted that this application is in condition for allowance and a notice to that effect is earnestly solicited. Should the Examiner have any questions regarding this response, or the application in general, he is invited to contact the undersigned at (540) 361-1863.

Respectfully submitted,

POTOMAC PATENT GROUP PLLC

By: /stevenmdubois/
Steven M. duBois
Registration No. 35,023

Date: December 10, 2008

Customer No. 42015
Potomac Patent Group PLLC
P.O. Box 270
Fredericksburg, VA 22404
(540) 361-1863